IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Application Serial No.	10/821,832
Filing Date	April 9, 2004
Inventorship	Rajiv K. Grover
Applicant/Appellant	Hewlett-Packard Company
Group Art Unit	2195
Examiner	WILSER, Michael P.
Confirmation No.	1395
Attorney's Docket No.	200402482-1

Title: Device Loading In Storage Networks

REPLY BRIEF

To: MS Appeal Brief - Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

As required, this Reply Brief is filed within two (2) months of the Examiner's Answer mailed November 12, 2009, and is in furtherance to the Appeal Brief.

This Reply Brief contains items under the following headings as required by 37 C.F.R. §41.37 and M.P.E.P §1206:

I. Status of Claims
II. Grounds of Rejection to be Reviewed on Appeal
III. Argument

I. STATUS OF CLAIMS

A. Total Number of Claims in Application

There are twenty-six (26) claims pending in this application (Claims 1-26).

B. Current Status of Claims

1. Claims canceled: none.
2. Claims withdrawn from consideration but not canceled: none.
3. Claims pending: 1-26.
4. Claims allowed: none.
5. Claims rejected: 1-26.

C. Claims on Appeal

The claims on appeal are claims 1-26.

II. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

The Office Action rejected claims 1-2, 4, 6-7, 9, 11-14, 16-21, 23-24, and 26 under 35 U.S.C. 103(a) as being unpatentable over U.S. Patent No. 7,225,242 to Cherian, et al ("Cherian") in view of the IBM publication to Tate, et al. ("Tate"). The Final Office Action also rejected claims 3, 5, 8, 10, 15, 22, and 25 under 35 U.S.C. 103(a) as being unpatentable over Cherian and Tate and further in view of U.S. Patent No. 2004/0078599 to Nahum ("Nahum"). Appellant requests the Board to review each of these grounds of rejection.

III. ARGUMENT

First Rejection under 35 U.S.C. §103(a)

Reply to Examiner's Answer (1)

Independent claim 1 recites "for each host port connection, determining actual loading of IO jobs for each of the storage devices based at least in part on a queue depth for each of the host port LUNs" (emphasis added).

The Examiner states that "Cherian teaches the server execution throttle is the maximum number of I/O commands that a server can have outstanding (col. 1, line 64 – col. 2, line 1), each LUN of the storage device is logically normally own [sic] by a single host server (col. 4, lines 56-68), and determine [sic] the execution throttle for each server(s) that has logical ownership over the LUN (col. 5, lines 8-10 and lines 13-15).

Appellant notes that that the commonly understood meaning of the word "actual" is "existing in act and not merely potentially." See, e.g., Merriam-Webster's online dictionary. Although not specifically called out for definition in Appellant's specification, this meaning is entirely consistent with use of the word "actual" throughout Appellant's specification and in the claims and prosecution history. For example, Appellant describes in para. [0030] that "[i]n an exemplary implementation, a loading factor may be applied so that the actual loading is always less than the maximum service queue depth." This is merely one example in the specification where Appellant clearly distinguishes that actual loading is different than the maximum service queue depth. Furthermore, the term "actual loading" is not to be confused with "device loading" which is used separately in Appellant's specification to refer to a not-to-exceed loading for a device. See, e.g., the equation in para. [0031].

Appellant does note an inconsistency in the specification at para. [0028] wherein the paragraph correctly states that device loading is being calculated, but incorrectly uses "actual loading" in the equation. Paragraph [0029] goes on to properly distinguish between "device loading" and "actual loading":

> *. . . Therefore, a storage network configured with the HPUX host and VA7100/7400 storage device, as described above, should function properly. In addition, this storage network may be designed with a device loading that is even higher, so long as the maximum actual IO loading is less than the acceptable device loading (e.g., 750).*

The Examiner relies on "server execution throttle" and "the maximum number of I/O commands" in Cherian as teaching "actual loading." However, "server execution throttle" and "the maximum number of I/O commands" in Cherian is at best akin to "maximum service queue depth" and "device loading" in Appellant's specification. None of these fit the definition of "existing in act and not merely potentially." To the contrary, "server execution throttle" and "the maximum number of I/O commands" in Cherian are merely existing potentially.

The Examiner also states that "the claim fails to define what and what can not be actual loading of I/O job. Bases [sic] on the specification (page 8, paragraph [0027] – page 9, paragraph [0029]), the actual loading of IO jobs is defined as maximum actual loading of the storage device . . . [t]herefore, Examiner interprets the execution throttle levels and potential command throughput of the servers is [sic] the same as the actual loading for each of the storage devices." Appellant notes that this interpretation of the reference is inconsistent with the commonly understood meaning of the word "actual." Indeed, the "potential command throughput" comment is directly opposed to the definition of "existing in act and not merely potentially."

The Examiner also noted that claim 1 is rejected based on Cherian and Tate, not Hinton and Tawil. Appellant agrees, and appreciates the Examiner noting this typographical error in the Appeal Brief.

Reply to Examiner's Answer (2)

Dependant claim 2 recites "determining actual loading for each of the storage devices based at least in part on a number of host groups in the storage network."

The Examiner refers to the discussion of "actual loading" in Examiner's Answer (1). Appellant has already addressed this above and therefore will not repeat the same argument for claim 2.

The Examiner then continues this same line of reasoning, stating "Cherian teaches determining actual loading for each of the storage devices based at least in part on a number of host groups in the storage network (Servers(A+B+C), Servers (B+C+D+E); col. 5, lines 23-27)." Appellant

The Examiner also states that "Appellant failed to provide any reasons/arguments why the execution throttle for various servers in Cherian is not the same as determining actual loading for each of the storage devices . . . ." Appellant has clarified above with regard to the arguments in support of claim 1. In addition, Appellant notes that a "throttle" is merely another way of saying "maximum" or "potential" throughput which is directly opposed to the definition of "actual (i.e., "existing in act and not merely potentially").

Reply to Examiner's Answer (3)

Claim 4 recites "uses a loading factor to determine if the actual loading for each of the storage devices exceeds a maximum loading."

The Examiner again relies on Equation 1 or Equation 2 as being the loading factor and states that Appellant "failed to provide any reasons/arguments why summing the execution throttle for various servers in Cherian is not the same as using a loading factor to determine if the actual loading for each of the storage devices exceeds a maximum loading."

With regard to these equations, Cherian states:

> *As to Equation 1, because each of Servers A, B, and C has logical ownership of a LUN on Storage Device X, the execution throttle of each of Servers A, B, and C is summed and compared to the command queue depth of Storage Device X. As to Equation 2, because each of Servers B, C, D, and E has logical ownership of a LUN on Storage Device Y, the execution throttle levels of each of Servers B, C, D, and E are summed and compared to the command queue depth of Storage Device Y.*

Summing the execution throttle for various servers in Cherian is not the same as using a loading factor. At most, these equations in Cherian are akin to Appellant's equations for determining maximum queue depth. See, e.g., equation in specification at para. [0027]. Appellant's specification goes on to discuss loading factors which may be applied after the maximum queue depth is determined. See, e.g., para. [0030] stating "a loading factor may be applied so that the actual loading is always less than the maximum service queue depth. For example, a loading factor of about 80-90% (e.g., maximum for user environments) or about 85-90% (e.g., minimum recommended for test environments) of the maximum service queue depth may be applied to one or more of the storage devices."

Reply to Examiner's Answer (4)

Claim 11 recites "A method providing an input/output (IO) flow control mechanism in a storage network." The Examiner provided no support for rejecting this recitation.

The Examiner states that "[a] preamble is generally not accorded any patentable weight where it merely recites the purpose of a process or the intended use of a structure, and where the body of the claim does not depend on the preamble for completeness but, instead, the process steps or structural limitations are able to stand alone." The Examiner appears to have accepted the "generally not" principal without further analysis.

If the Examiner wants to dismiss these recitations as being part of the preamble, it is error for the Examiner to summarily dismiss these recitations simply on the basis that these recitations are a part of the preamble. The MPEP states at Section 2111.02 (Effect of Preamble) that "[t]he determination of whether a preamble limits a claim is made on a case-by-case basis in light of the facts in each case; there is no litmus test defining when a preamble limits the scope of a claim."

The MPEP goes on in Section 2111.02 to state that "[t]he claim preamble must be read in the context of the entire claim. The determination of whether preamble recitations are structural limitations or mere statements of purpose or use "can be resolved only on review of the entirety of the [record] to gain an understanding of what the inventors actually invented and intended to encompass by the claim."

During examination, statements in the preamble reciting the purpose or intended use of the claimed invention must be evaluated to determine whether the recited purpose or intended use results in a structural difference between the claimed invention and the prior art.

There are two reasons why the recitations in claim 11 of "providing an input/output (IO) flow control mechanism in a storage network" recites a structural difference between the claimed invention and the prior art. First, providing an IO flow control mechanism is an affirmative act of providing a physical means to an end in a physical environment (i.e., the storage network). In fact, no function is recited by the act of providing an IO flow control mechanism in a storage network. That is, the result may be functional, but providing the means to achieve that result is not merely functional. Second, the Examiner was unable to find any support for this recitation in the cited references, providing evidence that this is indeed a distinction between the claimed invention and the prior art.

Therefore, it was error for the Examiner to simply dismiss these recitations during examination.

Reply to Examiner's Answer (5)

Claims 12-14 and 16-17 include further recitations for "automatically determining actual loading for the storage device."

The Examiner again relies on the execution throttle and Equation 1 in Cherian. Appellant notes that the execution throttle and Equation 1 do not teach determining the actual loading, for the reasons already set forth above. At most, Equation 1 sets forth a rule that may be checked, but there is no teaching of automatically determining an actual loading for the storage device.

Reply to Examiner's Answer (6)

Claim 18 recites "wherein the maximum loading for the storage device is based on a loading factor for test environments."

The Examiner states that "Cherian teaches 'the verification test is performed to determine whether the summed execution throttle value exceeds the command queue depth or command throughput of the associated controller' (col. 5, lines 17-21), and when the verification step is not satisfied for any storage controller, the execution throttle value for the server(s) is adjusted until the verification step is passed (col. 5, lines 38-67)."

Here, Cherian is merely stating that the execution throttle value is tested against the command queue depth or command throughput. There is no teaching of using a loading factor. Appellant describes in para. [0030] that "[i]n an exemplary implementation, a loading factor may be applied so that the actual loading is always less than the maximum service queue depth." Accordingly, the iterative test in Cherian is not needed because by applying a loading factor, the actual loading is always less than the maximum service queue depth.

Reply to Examiner's Answer (7)

Claim 19 recites "the loading factor is in the range of about 80% to 90% of the service queue depth for the storage device."

The Examiner summarily dismisses these recitations as being obvious "that the range should not be 100% of the service queue depth, and should not be too low, thus, 80-90% would be a safe choice for the system." There is no support in the references for either of these conclusions. Clearly this is nothing more than hindsight interpretation of the references in view of the claim recitations.

Appellant further notes that the Examiner stated in Examiner's Answer (6) above that "Cherian teaches 'the verification test is performed to determine whether the summed execution throttle value exceeds the command queue

depth or command throughput of the associated controller' (col. 5, lines 17-21), and when the verification step is not satisfied for any storage controller, the execution throttle value for the server(s) is adjusted until the verification step is passed (col. 5, lines 38-67)." This iterative test in Cherian is not needed if a loading factor is used because the actual loading is always less than the maximum service queue depth. Therefore, it is contrary to the teachings of Cherian to provide a loading factor.

Second Rejection under 35 U.S.C. §103(a)

Reply to Examiner's Answer (1)

Claim 3 recites "determining actual loading for each of the storage devices based at least in part on a number of LUN security groups in the storage network."

The Examiner previously admited that Cherian and Tate do not disclose these recitations. Instead, the Examiner relied on paragraph [0018] in Nahum. Appellant explained that paragraph [0018] in Nahum simply describes a security procedure for authenticating each host, but does not determine actual loading for each of the storage devices based at least in part on a number of LUN security groups in the storage network. Now, however, the Examiner states that "Cherian might qualified [sic] as LUN security group. The reference of Nahum is used to show LUN security groups are implemented in the storage device." Again, the Examiner relies on paragraph [0018] in Nahum.

Accordingly, the rejection is inconsistent and confusing at best, as to what exactly the Examiner is relying on in each of the cited references as being each of the claim recitations of "determining actual loading for each of the storage devices based at least in part on a number of LUN security groups in

the storage network." Appellant asserts that (for sake of argument) even if Nahum can be considered as teaching a LUN security group, there is no teaching of determining actual loading for each of the storage devices based at least in part on a number of LUN security groups in the storage network.

Reply to Examiner's Answer (2)

Claim 5 recites "the computer process further simplifies host groups and LUN security groups into virtual connections for analysis." Claim 10 includes similar recitations as claim 5.

The Examiner states that "Nahum teaches '[t]here is provided yet further a method for securing a host with at least one HBA which is identified by a first WWN (World Wide Number), and a storage device comprises at least one LU (logical unit), and identified by a second WWN and by at least one LUN (Logical Unit Number). The security procedure authenticates each one host out of the first array independently of the first WWN, and identifies each one storage device out of the second array by the second WWN and by the at least one LUN' (page 2, paragraph [0018])." This simply describes a process where one server and one storage device is authenticated at a time. It does not teach the computer process further simplifies host groups and LUN security groups into virtual connections for analysis.

Conclusion

For the reasons provided herein, Appellant respectfully requests the Board to rule that the rejections of the claims are improper.

Respectfully Submitted,

/Mark D. Trenner/

Dated: January 9, 2010 By: _____

Mark D. Trenner

Reg. No. 43,961

(720) 221-3708